SPC BENCHMARK 1™

FULL DISCLOSURE REPORT

HUARUI EXPON TECHNOLOGIES

EXPONTECH WDS V3

SPC-1™ V3.10.0

SUBMISSION IDENTIFIER: A32027

SUBMITTED FOR REVIEW: SEPTEMBER 19, 2023

REVISED – SECOND EDITION: 2023-10-13

Second Edition – September 2023

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESS OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by Huarui Expon Technologies for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

This publication was produced in China. Huarui Expon Technologies may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change with notice. Consult your local Huarui Expon Technologies representative for information on products and services available in your area.

Trademarks

SPC Benchmark 1, SPC-1, SPC-1 IOPS, SPC-1 LRT and SPC-1 Price-Performance are trademarks of the Storage Performance Council.

ExponTech and the ExponTech logo are trademarks or registered trademarks of Huarui Expon Technologies in China and other countries. All other brands, trademarks, and product names are the property of their respective owners.

Benchmark Specification and Glossary

The official SPC Benchmark 1™ (SPC-1™) specification is available on the website of the Storage Performance Council (SPC) at www.spcresults.org.

The SPC-1™ specification contains a glossary of the SPC-1™ terms used in this publication.

Table of Contents

AUDIT CERTIFICATION

InfoSizing
The Right Metric For Sizing IT

Storage Performance Council
Certified Auditor

Zong Chao
Huarui Expon Technologies
No.4011, Section A, Zhongguan Times Plaza
Liuxian Avenue, Taoyuan Street, Nanshan
District Shenzhen, Guangong CHINA

October 13, 2023

I verified the SPC Benchmark 1™ (SPC-1™ v3.10.0) test execution and performance results of the following Tested Storage Product:

ExponTech WDS V3

The results were:

SPC-1 IOPS™	**27,201,325**
SPC-1 Price-Performance	**$43.74/SPC-1 KIOPS™**
SPC-1 Total System Price	$1,189,577.99
SPC-1 IOPS Response Time	0.217 ms
SPC-1 Overall Response Time	0.144 ms
SPC-1 ASU Capacity	27,493 GB
SPC-1 ASU Price	$43.27/GB

In my opinion, these performance results were produced in compliance with the SPC requirements for the benchmark.

The testing was executed using the SPC-1 Toolkit Version v3.0.2. The audit process was conducted in accordance with the SPC Policies and met the requirements for the benchmark.

A Letter of Good Faith was issued by Huarui Expon Technologies, stating the accuracy and completeness of the documentation and testing data provided in support of the audit of this result.

Page 1 of 2
63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

A Full Disclosure Report for this result was prepared by InfoSizing, reviewed and approved by Huarui Expon Technologies, and can be found at www.spcresults.org under the Submission Identifier A32027.

The independent audit process conducted by InfoSizing included the verifications of the following items:

- The physical capacity of the data repository (76,816 GB).
- The total capacity of the Application Storage Unit (27,493 GB).
- The accuracy of the Benchmark Configuration diagram.
- The tuning parameters used to configure the Benchmark Configuration.
- The Workload Generator commands used to execute the testing.
- The validity and integrity of the test result files.
- The compliance of the results from each performance test.
- The compliance of the results from each persistence test.
- The compliance of the submitted pricing model.
- The differences between the tested and the priced configuration, if any.

The Full Disclosure Report for this result was prepared in accordance with the disclosure requirements set forth in the specification for the benchmark.

The following benchmark requirements, if any, were waived in accordance with the SPC Policies:

None.

Additional Notes:

This result was originally audited on September 19, 2023, with pricing in CNY. The pricing has been re-audited for publication using USD instead.

Respectfully Yours,

Doug Johnson, Certified SPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

LETTER OF GOOD FAITH

September 15,2023
To: Doug Johnson,SPC auditor
perflabs,Inc. DBA InfoSizing
63 Lourdes Drive
Leominster,MA 01453-6709
USA

Subject: SPC-1 Letter of Good Faith for the ExponTech WDS V3

Huarui Expon Technologies co., Ltd is the SPC-1 Test Sponsor for the above listed product. To the best of our knowledge and belief, the required SPC-1 benchmark results and materials we have submitted for that product are complete, accurate, and in full compliance with version 3.10 of the SPC-1 benchmark specification.

In addition, we have reported any items in the Benchmark Configuration and execution of the benchmark that affected the reported results even if the items are not explicitly required to be disclosed by the SPC-1 benchmark specification.

Sincerely,

GM of Storage Products Department
Huarui Expon Technologies Co., Ltd

2023, 9, 15
Date: September 15,2023

SPC Benchmark 1™

Executive Summary

ExponTech WDS V3

SPC-1 IOPS™	**27,201,325**
SPC-1 IOPS Response Time	**0.217 ms**
SPC-1 Overall Response Time	0.144 ms

SPC-1 Price Performance	**$43.74/SPC-1 KIOPS™**
SPC-1 Total System Price	**$1,189,577.99**
SPC-1 Overall Discount	65.08%
Currency / Target Country	USD / China
Availability Date	July 24, 2023

Extensions

SPC-1 Data Reduction	NA
SPC-1 Encryption	NA
SPC-1 NDU	NA
SPC-1 Synchronous Replication	NA
SPC-1 Snapshot	NA

Storage Metrics

SPC-1 Data Protection Level	Protected 2
SPC-1 Physical Storage Capacity	76,816 GB
SPC-1 ASU Capacity	27,493 GB
SPC-1 ASU Price	$43.27/GB

Priced Storage Configuration Summary

64	Mellanox CX624106AN-CDAT 100 Gb 2-port
1	ExponTech WDS V3
32	Storage Nodes
4,096	GB Total Cache
64	100 Gbps Total Front-End Ports
192	Total Storage Devices (400 GB NVMe)
2	H3C 100 Gbps Ethernet Switches
68	Total RUs

RAMPD Average Response Time (ms) vs. IOPS

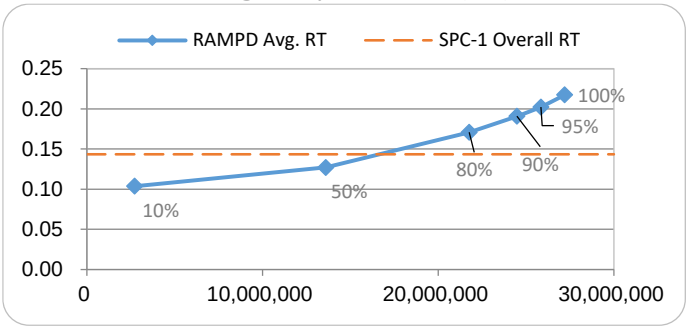

SUSTAIN Response Time (ms)

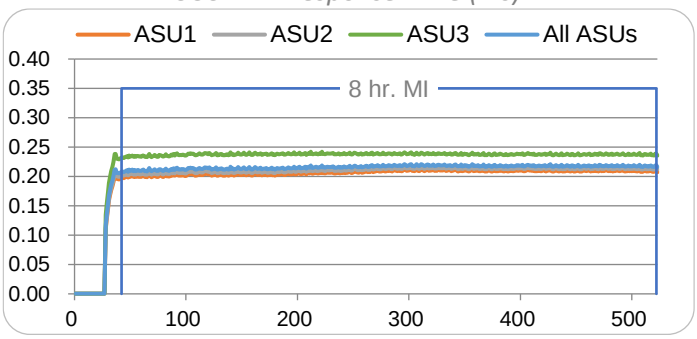

RAMPD_100 Response Time (ms)

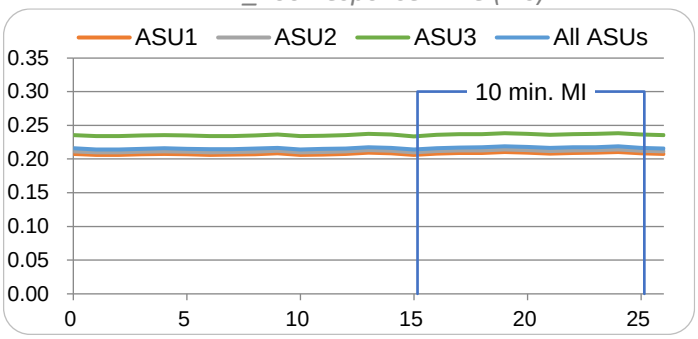

SPC Benchmark 1™ Specification Revision	v3.10.0
SPC Benchmark 1™ Workload Generator Revision	v3.0.2

Submitted for Review	September 19, 2023
Revised – Second Edition	October 13, 2023
Submission Details	www.storageperformance.org/r/A32027

PRICING DETAILS

Part No.	Description	Source	Qty	Unit Price	Ext. Price	Disc.	Disc. Price
Hardware & Software							
DST7250S	ExponTech WDS V3 Storage Node (2.5 inch *24) Case, Maximum expansion to 4096 Nodes in One System	1	32	5,698.63	182,356.16	65%	63,824.66
7250C53Y	Intel Xeon Gold 5318Y (2.1GHZ/24 cores) Processor Unit	1	64	6,575.34	420,821.76	65%	147,287.62
7250M32Y	Memory DDR4 RDIMM, 32GB 3200 MT/s	1	128	328.77	42,082.56	65%	14,728.90
7250100G	Network Card PCIEX16 100Gb Dual-port	1	96	3,068.49	294,575.04	65%	103,101.27
7250110C	Dapustor U.2 NVME SSD 400GB PCIe 4.0	1	192	3,835.62	736,439.04	65%	257,753.67
7250187C	Enterprise SSD 480GB SATA 2.5inch	1	64	493.15	31,561.60	65%	11,046.56
7250SSNS	ExponTech WDS V3 Basic Software, One Licese Matches One Node	1	32	32,876.71	1,052,054.72	65%	368,219.16
725010SW	H3C Ethernet Network Switch bundle 8 units* 100G*128 ports	1	2	87,671.23	175,342.46	65%	61,369.87
7250114G	10M 100G QSFP28 to QSFP28 AOC active optical cable	1	200	82.19	16,438.00	70%	4,931.40
7250145G	Mellanox Optical Module SFP 100G MM	1	400	876.71	350,684.00	70%	105,205.20
7250210G	Storage RACK 42U AC Cabinet	1	2	1,095.89	2,191.78	60%	876.72
					Hardware & Software Subtotal		**1,138,345.03**
Support & Maintenance							
7250201S	Installation Service - Engineering	1	32	410.96	13,150.72	50%	6,575.36
7250206U	Upgrade to Onsite Premier 24x7x4H Engineer Onsite Service - 36Month(s)	1	32	2,739.73	87,671.36	50%	43,835.68
7250206S	Switch Bundle Installation Service Engineering	1	2	821.92	1,643.84	50%	821.92
					Support & Maintenance Subtotal		**51,232.96**
SPC-1 Total System Price							**1,189,577.99**
SPC-1 IOPS™							27,201,325
SPC-1 Price-Performance™ ($/SPC-1 kIOPS™)							**43.74**
SPC-1 ASU Capacity (GB)							27,493
SPC-1 ASU Price ($/GB)							**43.27**

Discount Details: The discounts shown are based on the storage capacity purchased and are generally available.

Warranty: The priced maintenance provides 7x24x4H arrival service within designated city and distance. The service includes 7x24 contact to the Expontech call center with 4-hours on-site hardware replacement or troubleshooting, and online software support with access to all new software updates or troubleshooting.]

Differences Between Tested and Priced Storage Configurations

There were no differences between the TSC and the Priced Storage Configuration.

PUBLICATION DETAILS

This section provides contact information for the test sponsor and auditor, a revision history of this document, and a description of any exceptions or waivers associated with this publication.

Contact Information

Role	Name	Details
Test Sponsor Primary Contact	Huarui Expon Technologies Zong Chao	www.expontech.com zongchao@expontech.com
SPC Auditor	InfoSizing Doug Johnson	www.sizing.com doug@sizing.com

Revision Information

Date	FDR Revision	Details
September 19, 2023	First Edition	Initial Publication
October 13, 2023	Second Edition	Convert pricing from CNY to USD

Anomalies, Exceptions, Waivers

There were no anomalies, exceptions or waivers associated with the audit of the ExponTech WDS V3.

CONFIGURATION INFORMATION

Tested Storage Product Description

Expontech WDS V3 is a fully self-developed software-defined, high-performance distributed enterprise-level block storage platform designed for large-scale core data processing applications. ExponTech WDS V3 adopts a decentralized distributed system architecture, enabling smooth horizontal expansion while maintaining linear scalability of performance and capacity with the number of nodes. The platform supports NVMe SSD, SATA SSD, and SCM drives as primary storage media and is compatible with both 25G/100G RDMA RoCEv2 and traditional 10G TCP/IP networking technologies. With its advantages of high performance, scalability, reliability, ease of management, and maintenance, ExponTech WDS V3 fully meets the storage and disaster recovery requirements of massive data in fields such as cloud computing, big data, virtualization, databases, and high-performance computing. It has been widely used in industries such as finance, healthcare, education, government, enterprise, telecommunications, and transportation.

Host System and Tested Storage Configuration Components

The following table lists the components of the Host System(s) and the TSC.

Host Systems
32x H3C UniServer R4900 G5 Servers, each with : 2x Intel Xeon Gold 5318Y 2.1 GHz 24-Core Processor 256 GB Main Memory CentOS Linux release 7.9.2009
Tested Storage Configuration
#64x Mellanox CX623106AN-CDAT 100 Gbps 2-port HBAs
1x ExponTech WDS V3 with: **32x ExponTech ET-WDS storage node, each with:** **128 GB cache (4,096 GB total)** **2x 100 Gbps Front End Ports (64 total)** **192x 400 Gb NVME Storage Devices**
2x H3C S9820-8C 100 Gbps Ethernet Switch (128 active ports)

Component Changes in Revised Full Disclosure Report

The following table outlines component changes that were made in revisions to this Full Disclosure Report.

Original Component	Revised Component	Description of Change
n/a	n/a	Initial submission

Configuration Diagrams

BC/TSC Configuration Diagram

The following diagram illustrates the Benchmark Configuration (BC), including the Tested Storage Configuration (TSC) and the Host System(s).

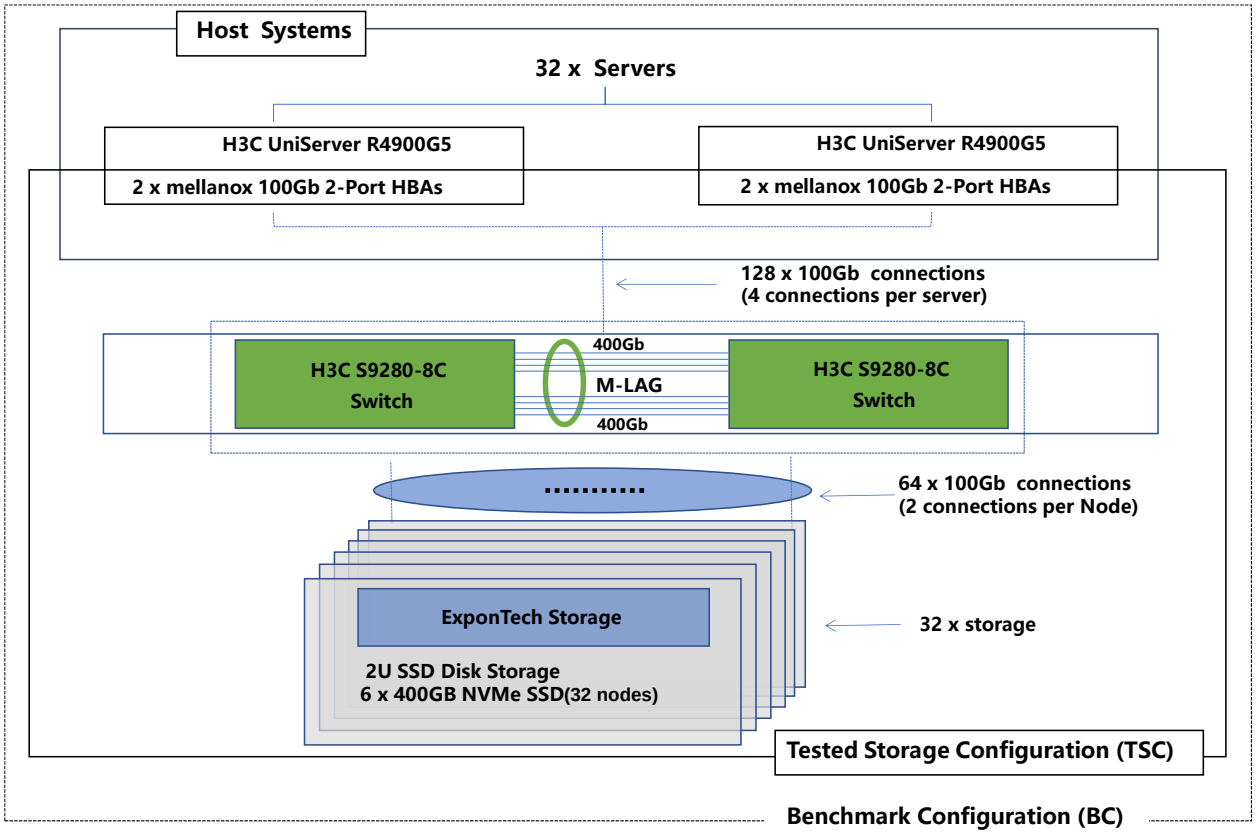

Storage Network Configuration

The Tested Storage Configuration (TSC) comprised 32 ExponTech WDS V3 storage nodes driven by 32 host systems (H3C UniServer R4900 G5). Each host system had two 100 Gb Mellanox MCX623106AN-CDAT connections to 2 switches (H3C S9820-8C) respectively. This was a total of 128 100Gb connections between the hosts and the switches. Each ExponTech WDS V3 storage node had 1 Mellanox MCX623106AN-CDAT (2 connections) to each H3C S9820-8C switch. This is a total of 64 100 Gb connections between the storage nodes and the switches.

Benchmark Configuration Creation Process

Customer Tuning Parameters and Options

All the customer tuning parameters and options that have been altered from their default values for this benchmark are included in Appendix C and in the Supporting Files (see Appendix A).

Tested Storage Configuration Creation

A detailed description of how the logical representation of the TSC was created is included in Appendix D and in the Supporting Files (see Appendix A).

Tested Storage Configuration Inventory

An inventory of the components in the TSC, as seen by the Benchmark Configuration, is included in Appendix E and in the Supporting Files (see Appendix A).

Workload Generator Storage Configuration

The SPC-1 Workload Generator storage configuration commands and parameters used to invoke the execution of the tests are included in Appendix F and in the Supporting Files (see Appendix A).

Logical Volume Capacity and Application Storage Unit Mapping

The following table details the capacity of the Application Storage Units (ASUs) and how they are mapped to logical volumes (LVs). All capacities are reported in GB.

	LV per ASU	LV Capacity	Used per LV	Total per ASU	% ASU Capacity	Optimized*
ASU-1	18	687.3	687.3	12,371.9	45.0%	No
ASU-2	18	687.3	687.3	12,371.9	45.0%	No
ASU-3	4	687.3	687.3	2,749.3	10.0%	No
	SPC-1 ASU Capacity			27,493		

*See Space Optimization Techniques

Physical Storage Capacity and Utilization

The following table details the Physical Capacity of the storage devices and the Physical Capacity Utilization (percentage of Total Physical Capacity used) in support of hosting the ASUs. All capacities are reported in GB.

Devices	Count	Physical Capacity	Total Capacity
400 GB NVMe	192	400.1	76,817.0
	Total Physical Capacity		**76,816**
	Physical Capacity Utilization		**35.79%**

Data Protection

The data protection level used for all LVs was **Protected 2 (Replication)**, which was accomplished providing fully redundant pathways from each host to the storage cluster where all data was replicated and distributed on two separate storage nodes.

Space Optimization Information

Description of Utilized Techniques

The TSC did not use any space optimization techniques.

Physical Free Space Metrics

The following table lists the Physical Free Space as measured at each of the required points during test execution. If space optimization techniques were not used, "NA" is reported.

Physical Free Space Measurement	Free Space (GB)
After Logical Volume Creation	NA
After ASU Pre-Fill	NA
After Repeatability Test Phase	NA

Space Optimization Metrics

The following table lists the required space optimization metrics. If space optimization techniques were not used, "NA" is reported.

Metric	Value
SPC-1 Space Optimization Ratio	NA
SPC-1 Space Effectiveness Ratio	NA

BENCHMARK EXECUTION RESULTS

This portion of the Full Disclosure Report documents the results of the various SPC-1 Tests, Test Phases, and Test Runs.

Benchmark Execution Overview

Workload Generator Input Parameters

The SPC-1 Workload Generator commands and input parameters for the Test Phases are presented in the Supporting Files (see Appendix A).

Measurement Intervals by Test Phase Graph

The following graph presents the average IOPS and the average Response Times measured over the MI of each Test Phase.

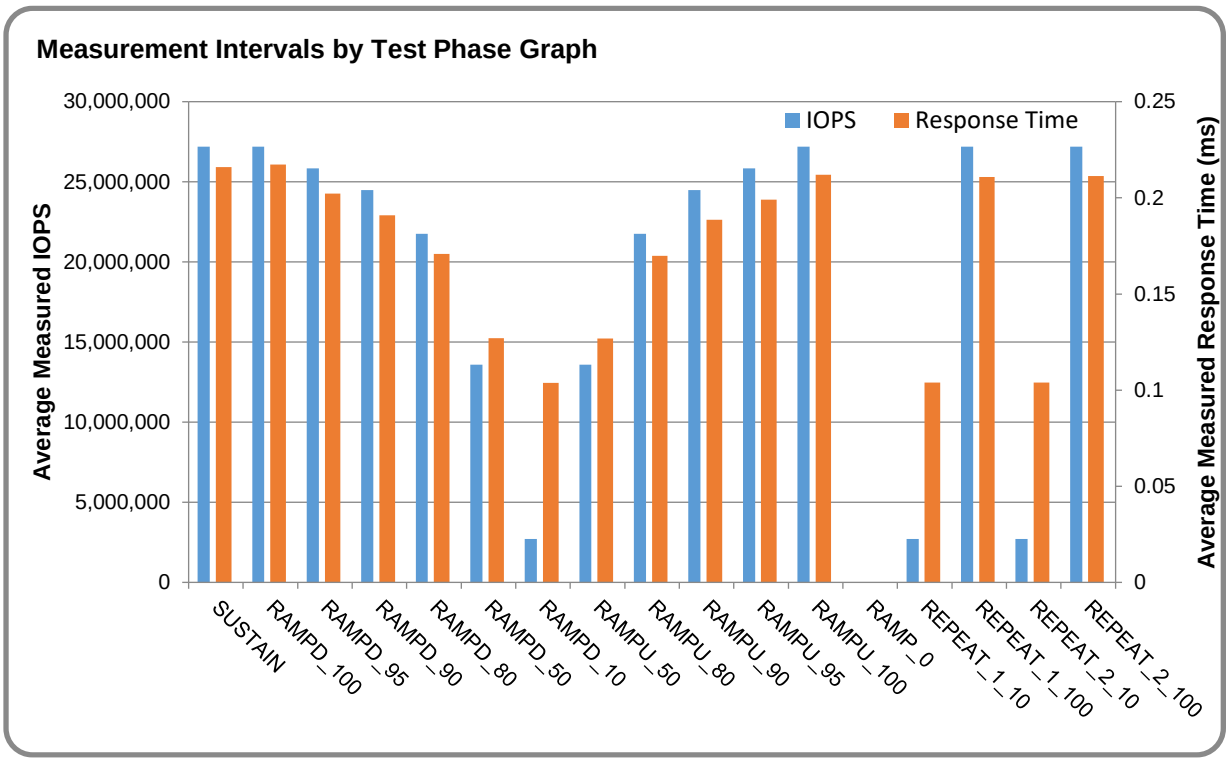

Response Time vs. Throughput Graph

The following graph presents the average Response Times versus the average IOPS for RAMPD_100 to RAMPD_10.

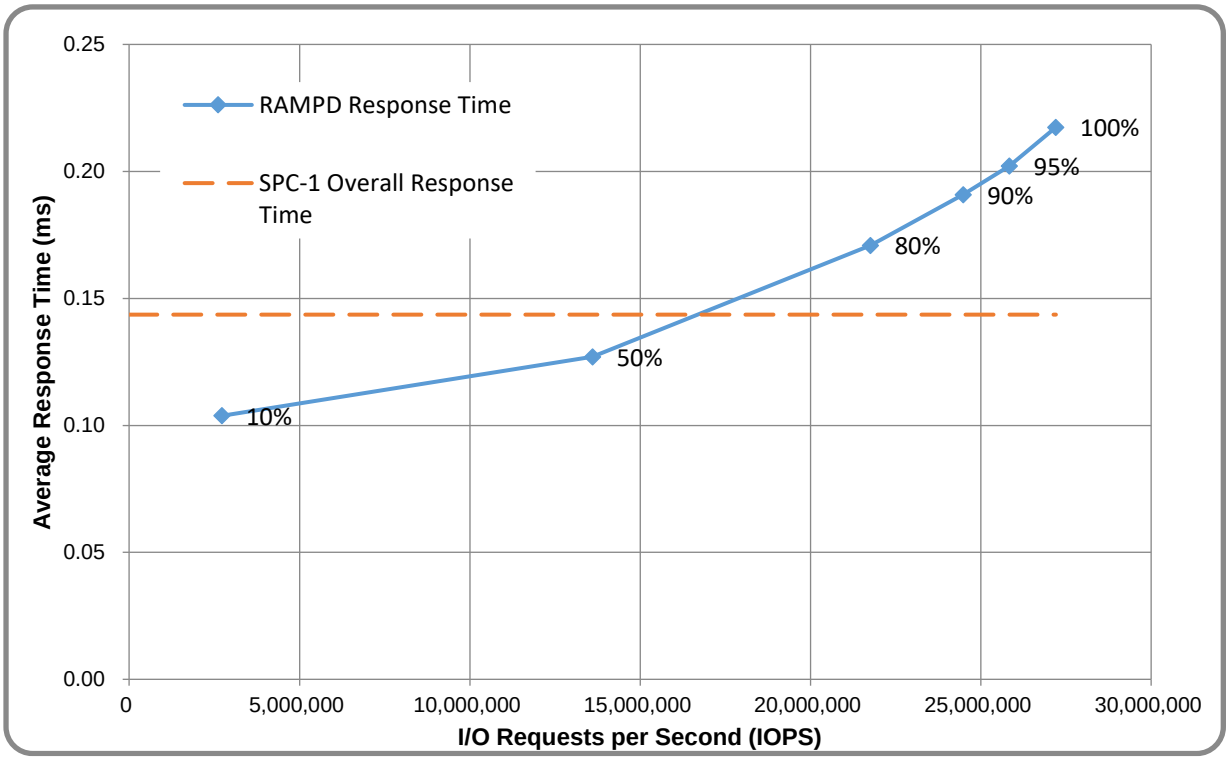

ASU Pre-Fill

The following table provides a summary of the Pre-Fill performed on the ASU prior to testing.

ASU Pre-Fill Summary			
Start Time	31-Aug-23 17:38:50	**Requested IOP Level**	100,000 MB/sec
End Time	31-Aug-23 17:47:43	**Observed IOP Level**	51,573 MB/sec
Duration	0:08:53	For additional details see the Supporting Files.	

SUSTAIN Test Phase

SUSTAIN – Results File

The results file generated during the execution of the SUSTAIN Test Phase is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_METRICS_0_Raw_Results.xlsx**

SUSTAIN – Execution Times

Interval	Start Date & Time	End Date & Time	Duration
Transition Period	31-Aug-23 18:21:05	31-Aug-23 18:36:05	0:15:00
Measurement Interval	31-Aug-23 18:36:05	01-Sep-23 02:36:06	8:00:01

SUSTAIN – Throughput Graph

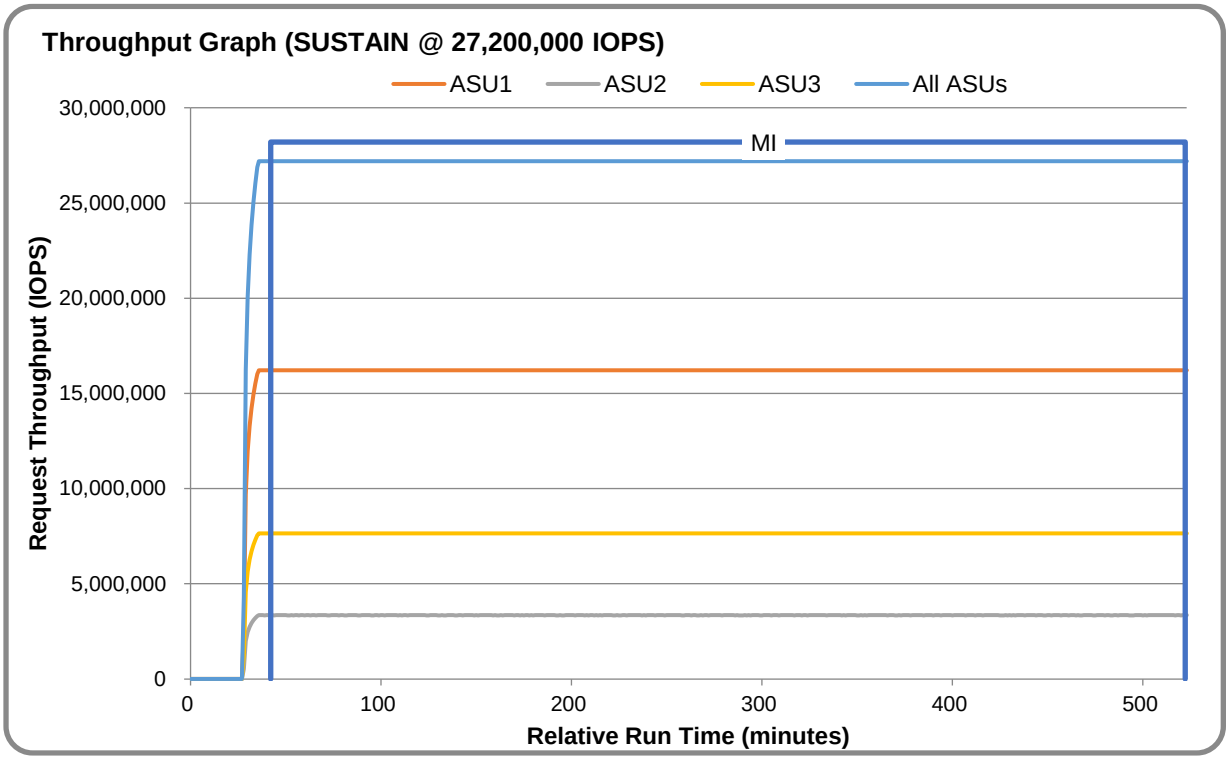

SUSTAIN – Response Time Graph

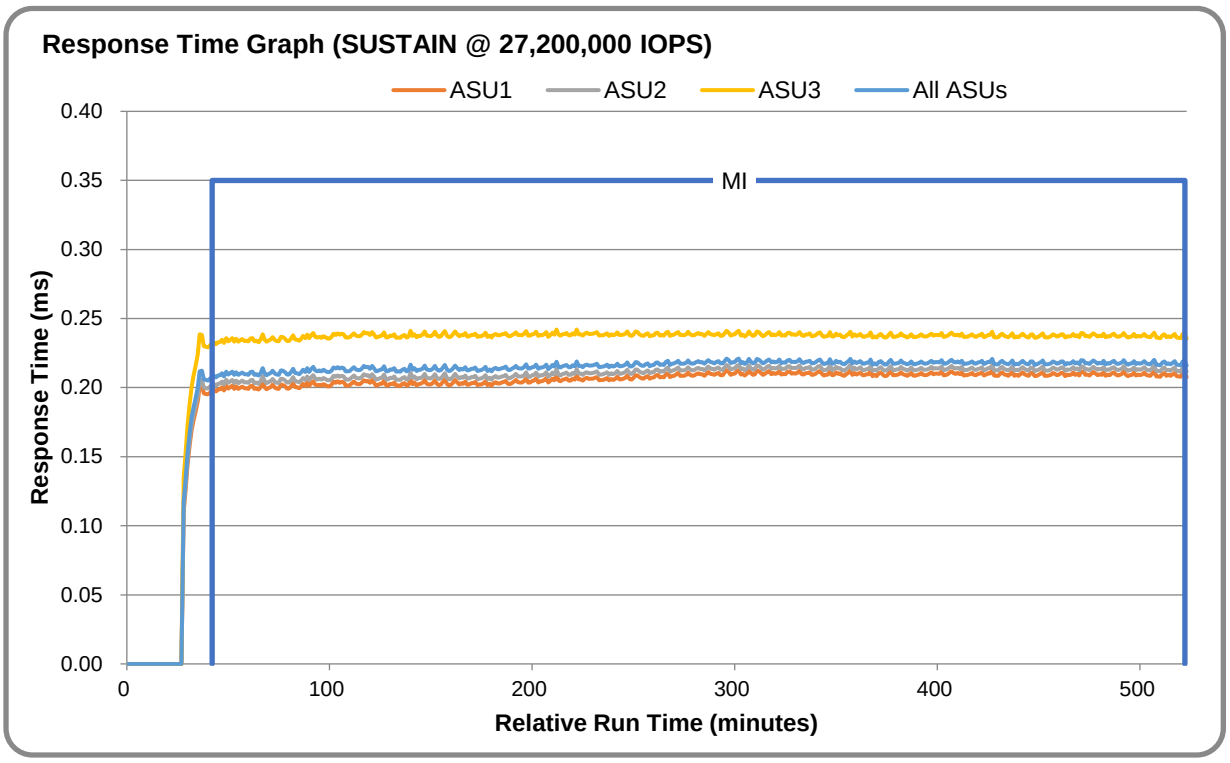

SUSTAIN – Data Rate Graph

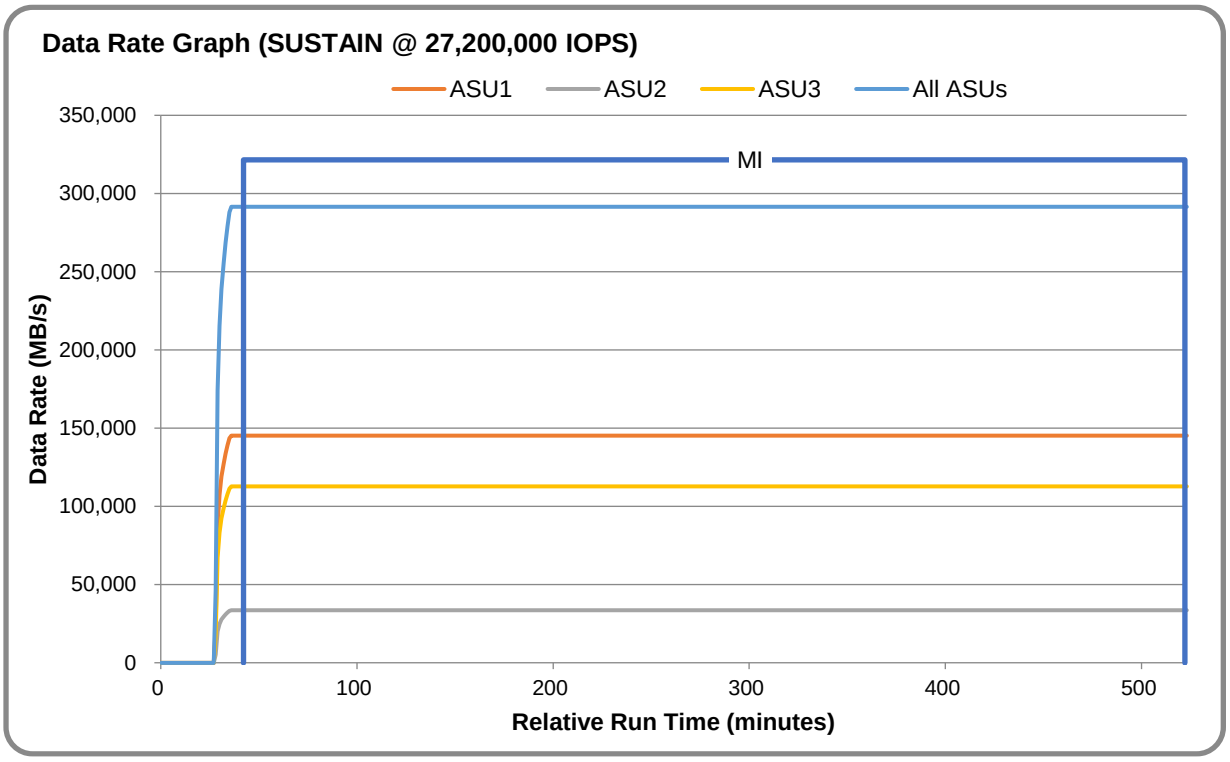

SUSTAIN – Response Time Frequency Graph

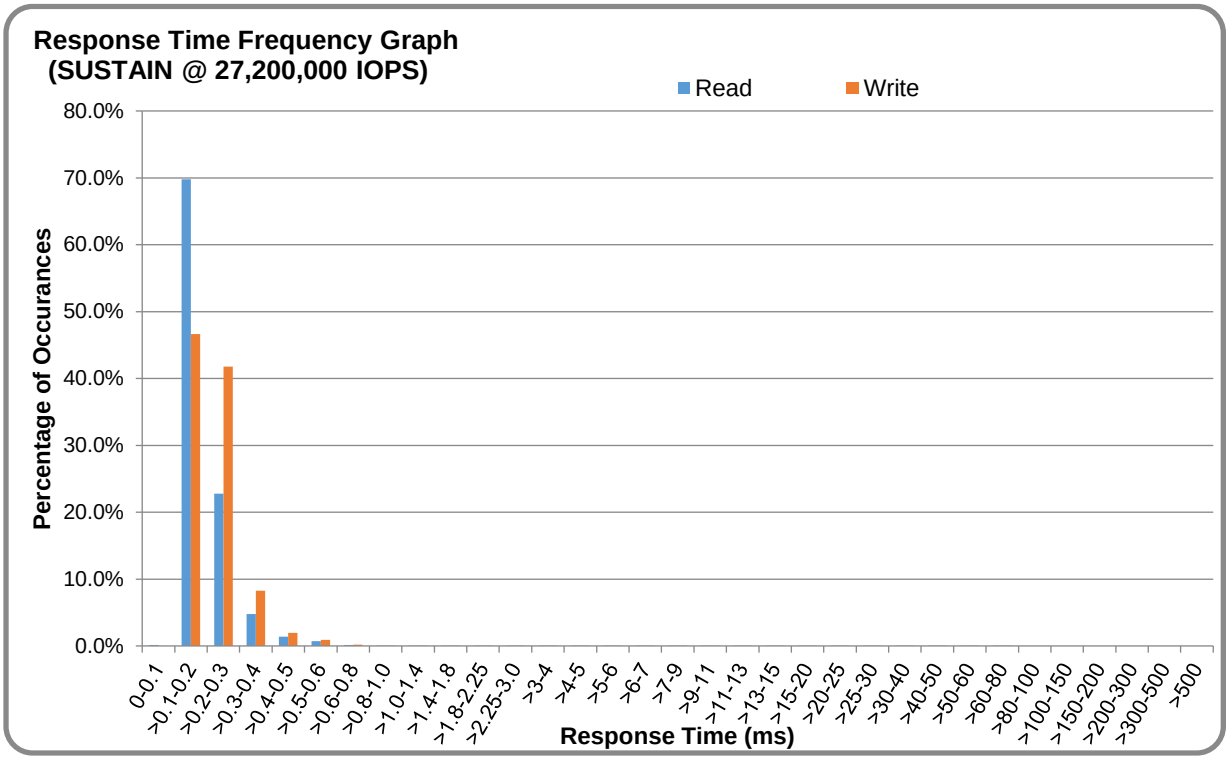

SUSTAIN – Intensity Multiplier

The following table lists the targeted intensity multiplier (Defined), the measured intensity multiplier (Measured) for each I/O stream, its coefficient of variation (Variation), and the percentage of difference (Difference) between Defined and Measured.

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0001	0.0000	0.0001	0.0000	0.0002	0.0001	0.0001	0.0000
Difference	0.004%	0.002%	0.003%	0.000%	0.006%	0.004%	0.005%	0.002%

RAMPD_100 Test Phase

RAMPD_100 – Results File

The results file generated during the execution of the RAMPD_100 Test Phase is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_METRICS_0_Raw_Results.xlsx**

RAMPD_100 – Execution Times

Interval	Start Date & Time	End Date & Time	Duration
Transition Period	01-Sep-23 02:37:06	01-Sep-23 02:52:07	0:15:01
Measurement Interval	01-Sep-23 02:52:07	01-Sep-23 03:02:07	0:10:00

RAMPD_100 – Throughput Graph

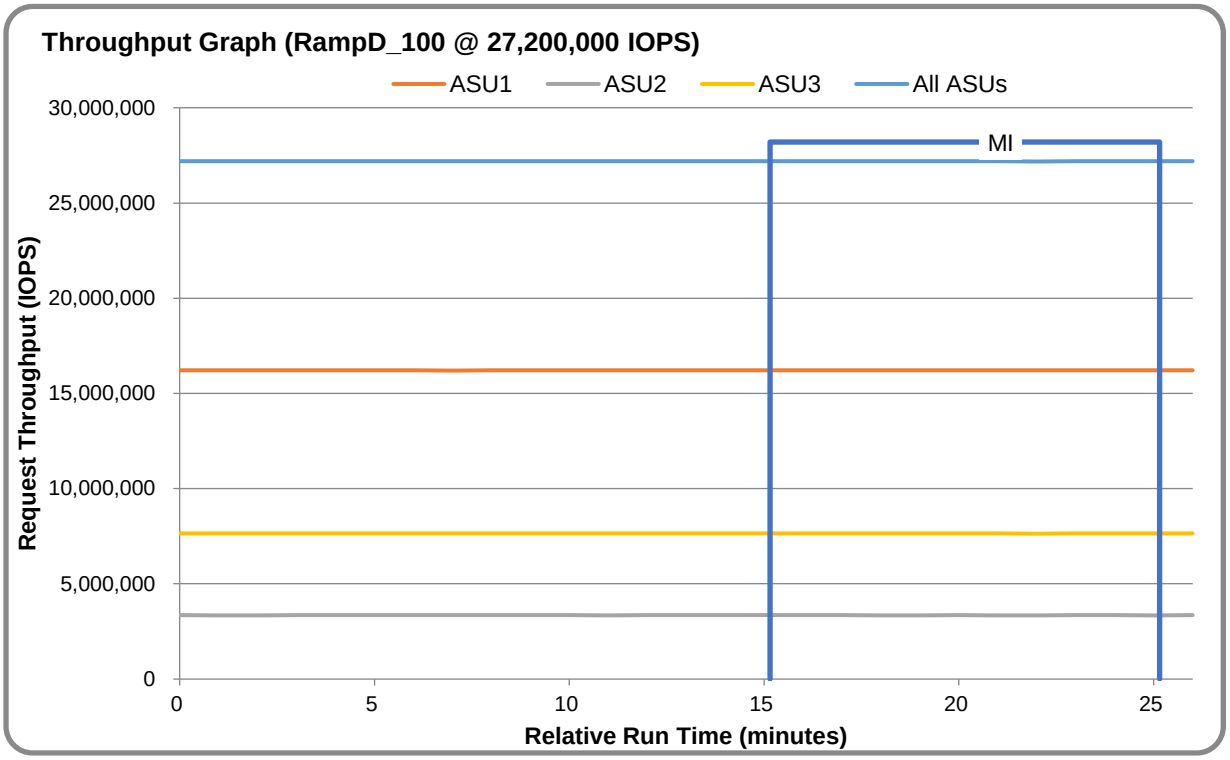

RAMPD_100 – Response Time Graph

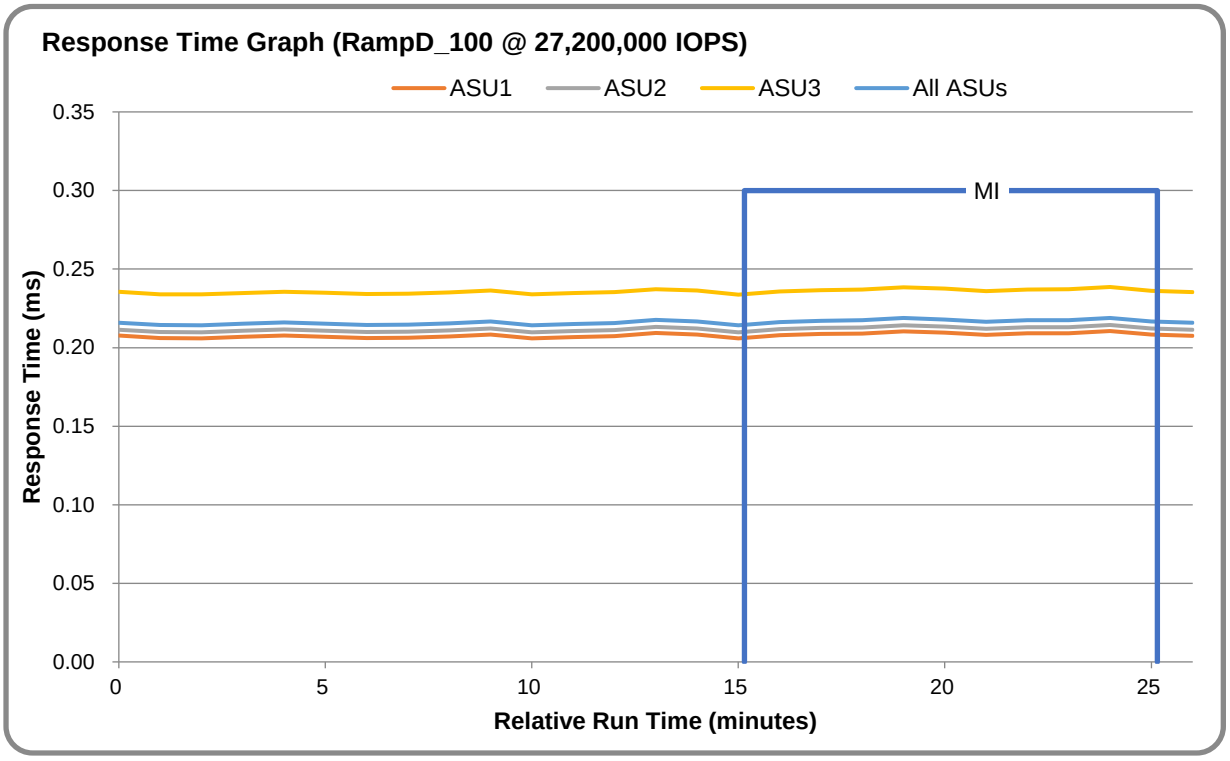

RAMPD_100 – Data Rate Graph

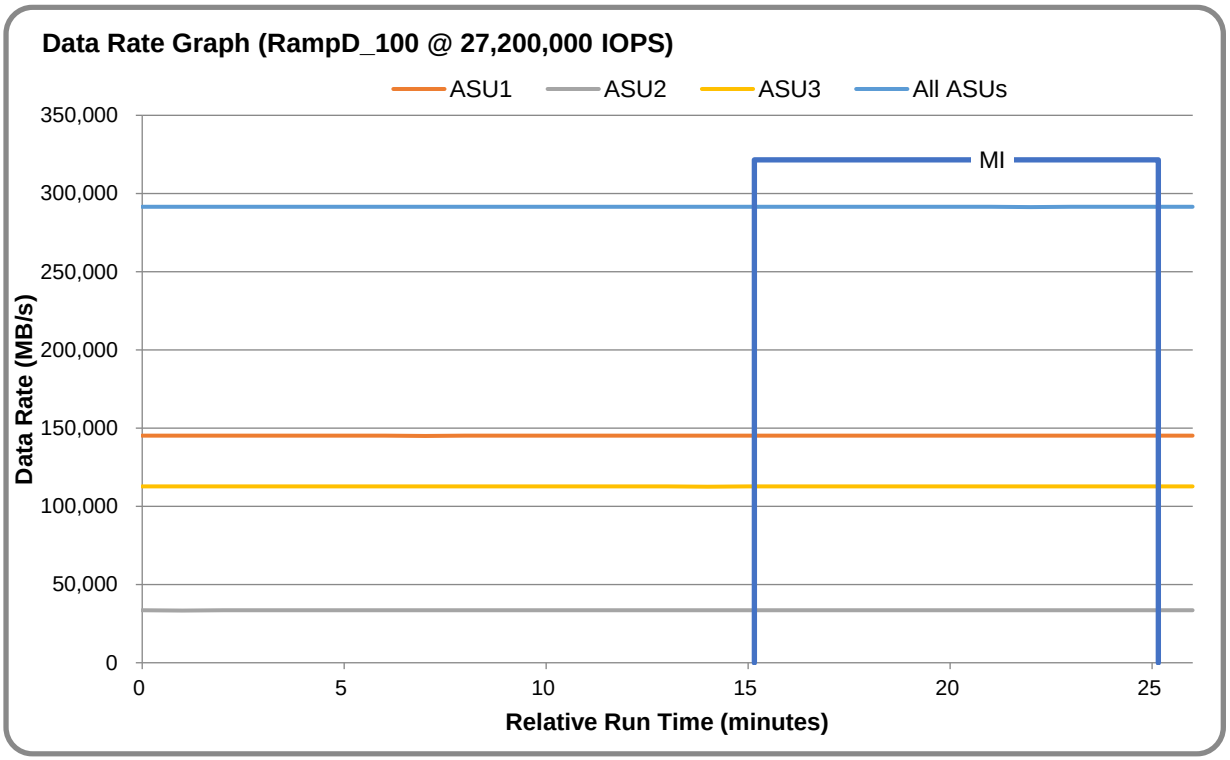

RAMPD_100 – Response Time Frequency Graph

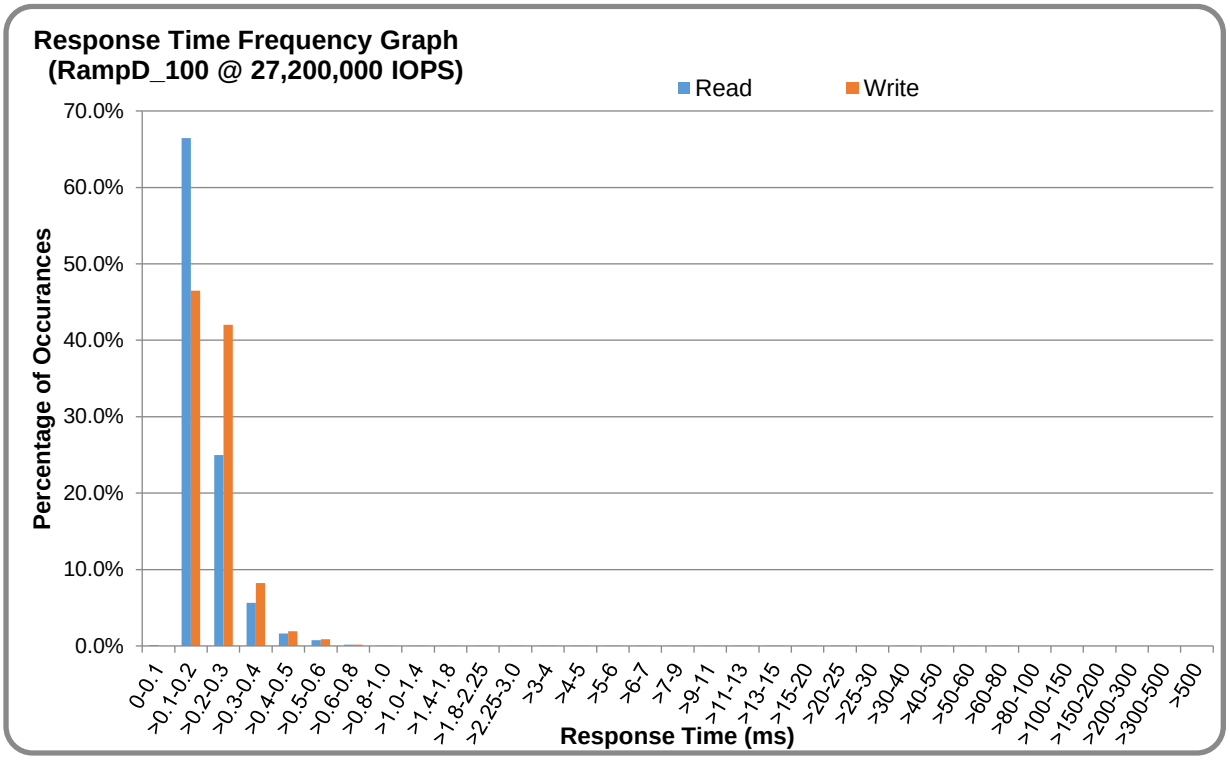

RAMPD_100 – Intensity Multiplier

The following table lists the targeted intensity multiplier (Defined), the measured intensity multiplier (Measured) for each I/O stream, its coefficient of variation (Variation), and the percentage of difference (Difference) between Defined and Measured.

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0002	0.0000	0.0001	0.0001	0.0003	0.0001	0.0001	0.0001
Difference	0.000%	0.002%	0.003%	0.001%	0.012%	0.005%	0.006%	0.002%

RAMPD_100 – I/O Request Summary

I/O Requests Completed in the Measurement Interval	16,320,829,797
I/O Requests Completed with Response Time <= 30 ms	16,320,783,225
I/O Requests Completed with Response Time > 30 ms	46,572

Response Time Ramp Test

Response Time Ramp Test – Results File

The results file generated during the execution of the Response Time Ramp Test is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_METRICS_0_Raw_Results.xlsx**

Response Time Ramp Test – Phases

The Response Time Ramp Test is comprised of 11 Test Phases, including six Ramp-Down Phases (executed at 100%, 95%, 90%, 80%, 50%, and 10% of the Business Scaling Unit) and five Ramp-Up Phases (executed at 50%, 80%, 90%, 95%, and 100% of the Business Scaling Unit).

Response Time Ramp Test – Average Throughput Graph

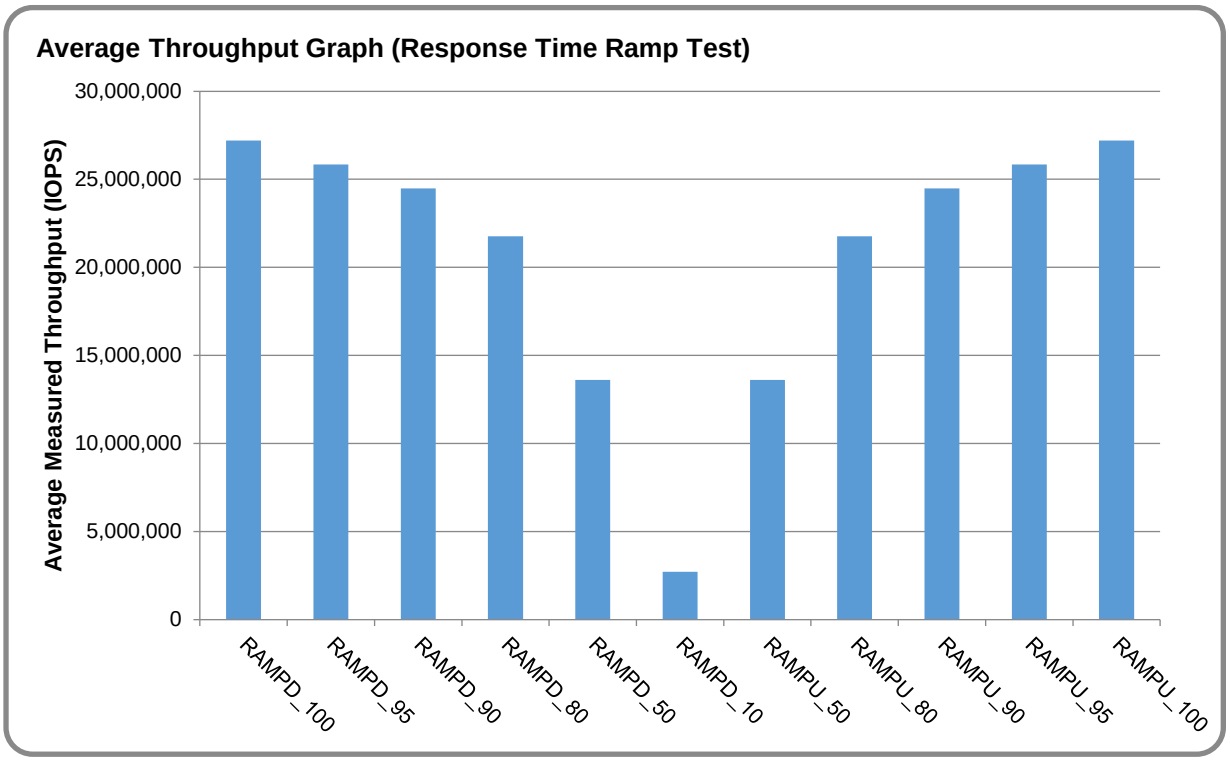

Response Time Ramp Test – Average Response Time Graph

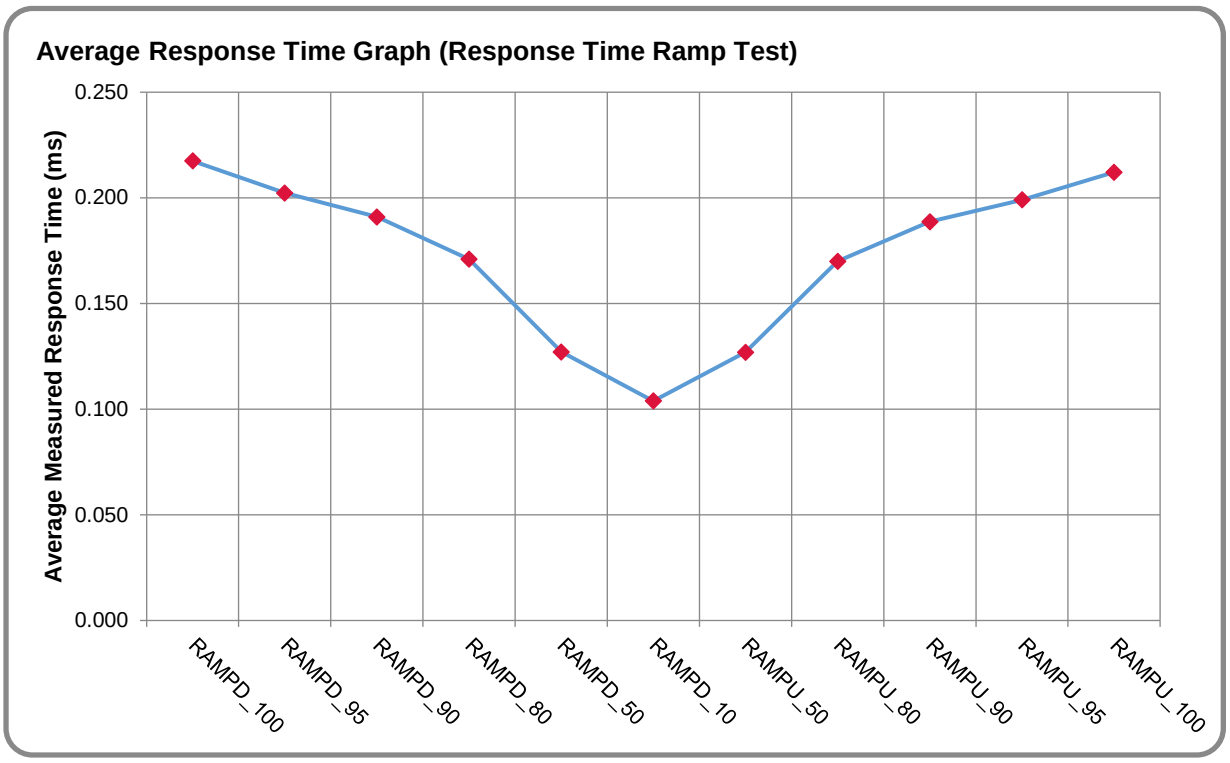

Response Time Ramp Test – RAMPD_10 Response Time Graph

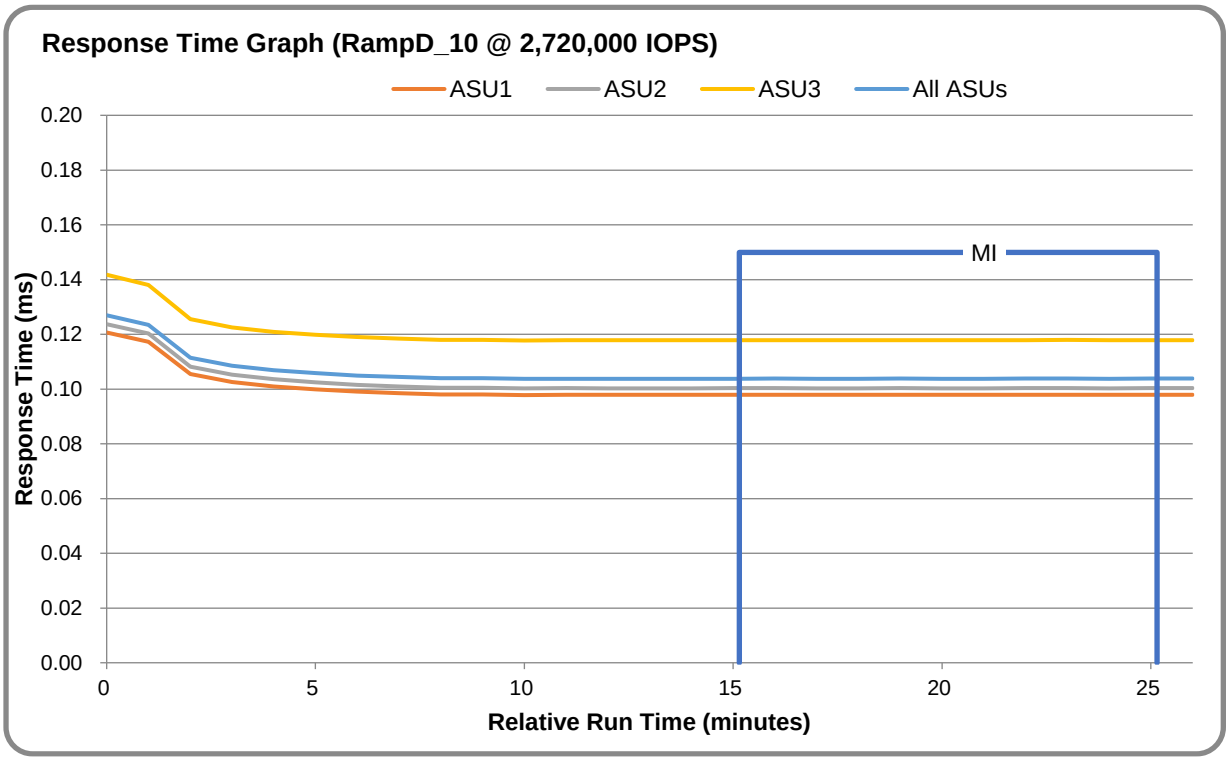

Repeatability Test

Repeatability Test Results File

The results file generated during the execution of the Repeatability Test is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_METRICS_0_Raw_Results.xlsx**

Repeatability Test Results

The throughput measurements for the Response Time Ramp Test (RAMPD) and the Repeatability Test Phases (REPEAT_1 and REPEAT_2) are listed in the table below.

Test Phase	100% IOPS	10% IOPS
RAMPD	27,201,325.6	2,720,131.1
REPEAT_1	27,201,971.9	2,720,149.4
REPEAT_2	27,201,625.5	2,720,207.1

REPEAT_1_100 – Throughput Graph

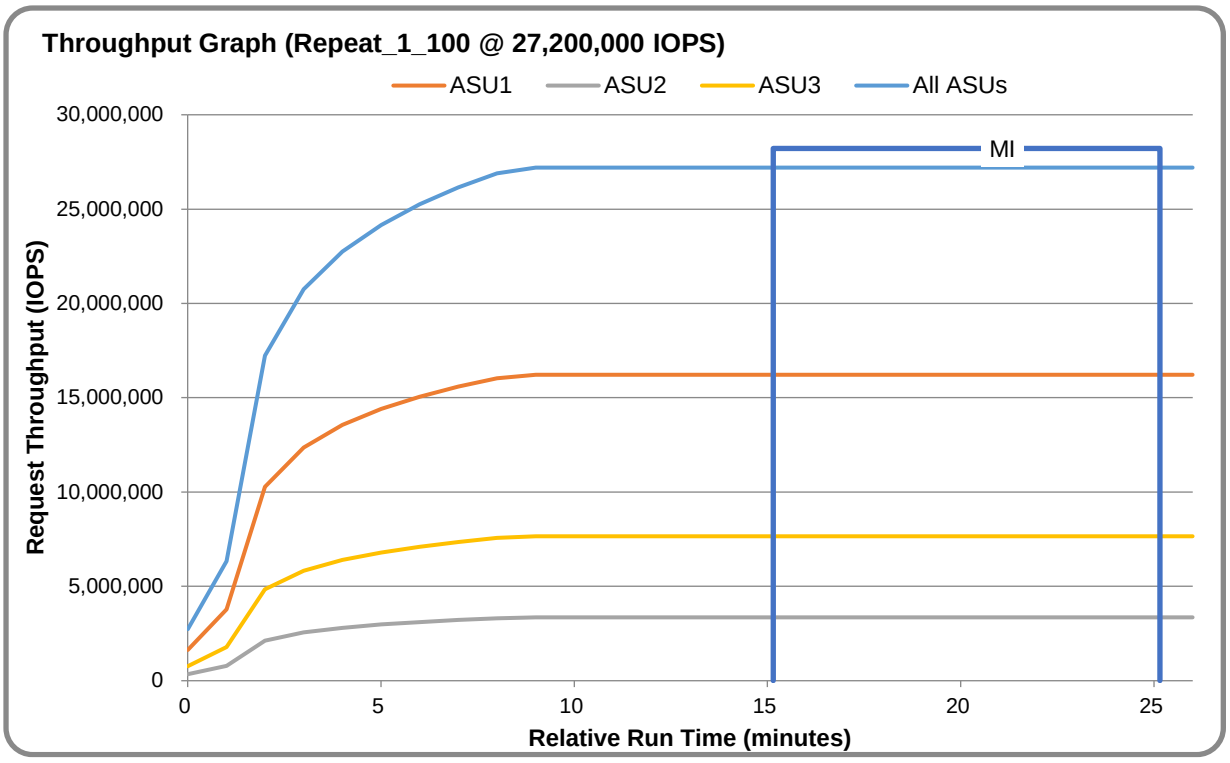

REPEAT_1_100 – Response Time Graph

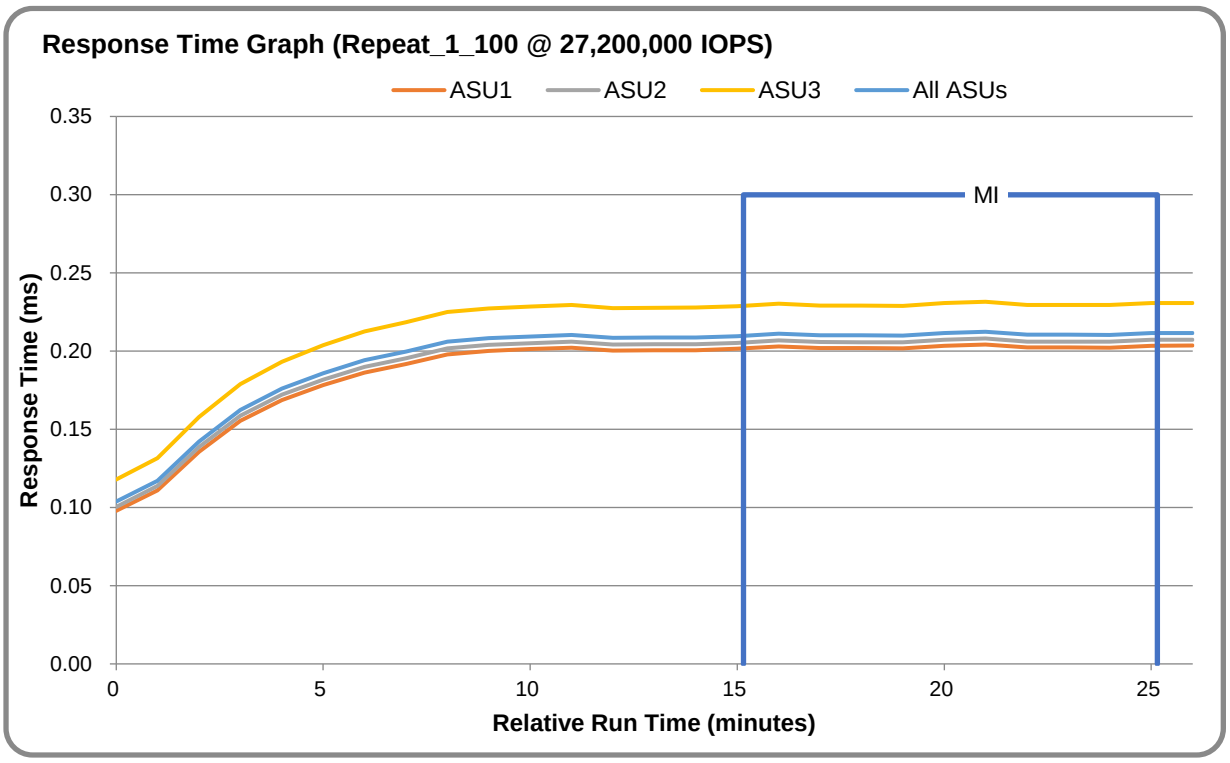

REPEAT_2_100 – Throughput Graph

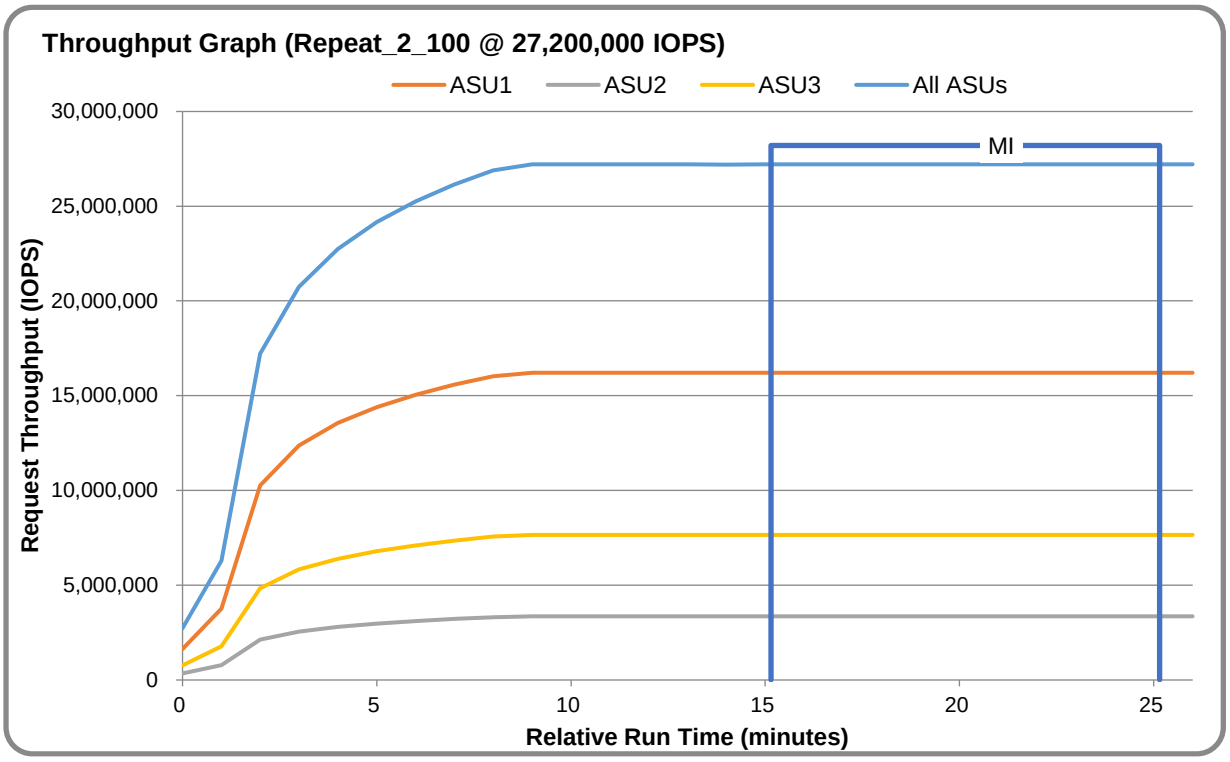

REPEAT_2_100 – Response Time Graph

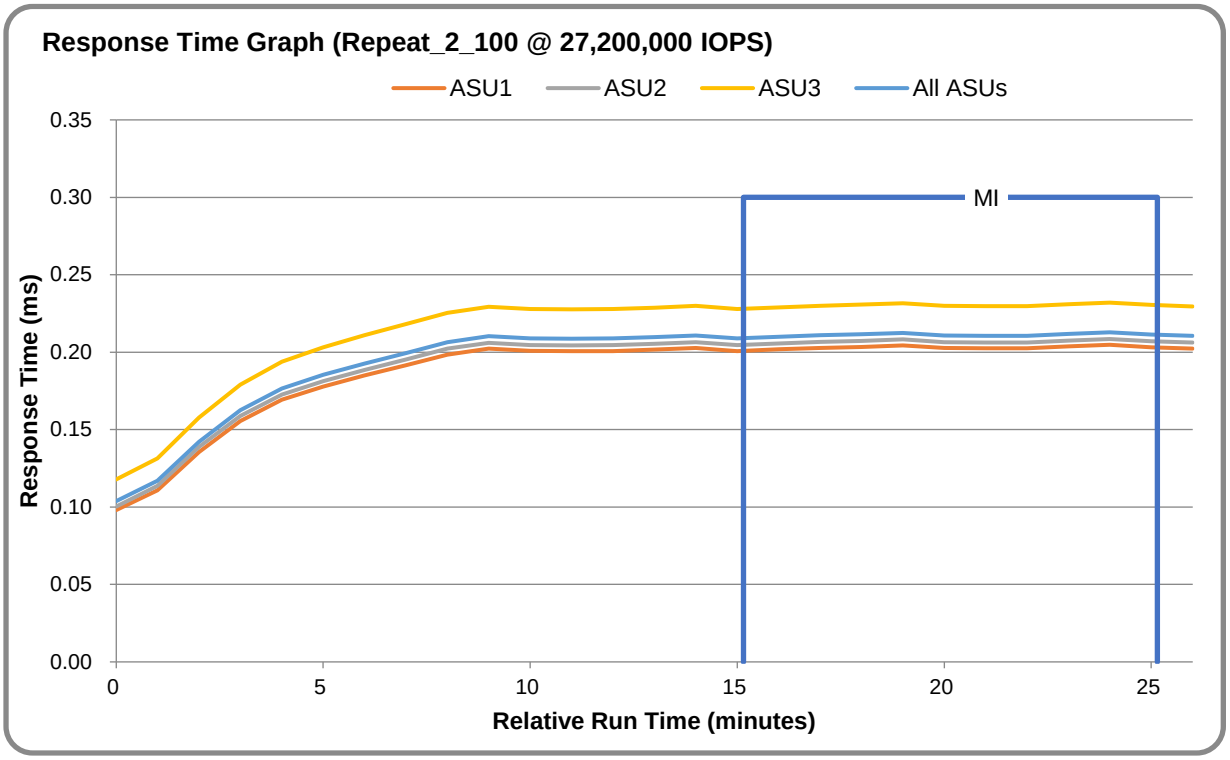

Repeatability Test – Intensity Multiplier

The following tables lists the targeted intensity multiplier (Defined), the measured intensity multiplier (Measured) for each I/O stream, its coefficient of variation (Variation), and the percent of difference (Difference) between Defined and Measured.

REPEAT_1_100 Test Phase

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0002	0.0000	0.0001	0.0000	0.0003	0.0001	0.0001	0.0000
Difference	0.007%	0.006%	0.006%	0.002%	0.005%	0.007%	0.005%	0.001%

REPEAT_2_100 Test Phase

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0001	0.0001	0.0001	0.0000	0.0001	0.0001	0.0002	0.0001
Difference	0.006%	0.005%	0.007%	0.000%	0.007%	0.003%	0.005%	0.000%

Data Persistence Test

Data Persistence Test Results File

The results files generated during the execution of the Data Persistence Test is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_PERSIST_1_0_Raw_Results.xlsx**
- **SPC1_PERSIST_2_0_Raw_Results.xlsx**

Data Persistence Test Execution

The Data Persistence Test was executed using the following sequence of steps:

- The PERSIST_1_0 Test Phase was executed to completion.
- The Benchmark Configuration was taken through an orderly shutdown process and powered off.
- The Benchmark Configuration was powered on and taken through an orderly startup process.
- The PERSIST_2_0 Test Phase was executed to completion.

Data Persistence Test Results

Data Persistence Test Phase: Persist1	
Total Number of Logical Blocks Written	1,341,418,014
Total Number of Logical Blocks Verified	556,335,860
Total Number of Logical Blocks Overwritten	785,082,154
Total Number of Logical Blocks that Failed Verification	0
Time Duration for Writing Test Logical Blocks (sec.)	601
Size in bytes of each Logical Block	8,192
Number of Failed I/O Requests in the process of the Test	0

Committed Data Persistence Implementation

ExponTech WDS V3 uses NVME SSD as the storage carrier for data and metadata. When data is written to the storage cluster, I/O will not return success until all data (including replicated data and meta data) are written to the NVME SSD disk.

APPENDIX A: SUPPORTING FILES

The following table details the content of the Supporting Files provided as part of this Full Disclosure Report.

File Name	Description	Location
/SPC1_RESULTS	**Data reduction worksheets**	**root**
SPC1_INIT_0_Raw_Results.xlsx	Raw results for INIT Test Phase	/SPC1_RESULTS
SPC1_METRICS_0_Quick_Look.xlsx	Quick Look Test Run Overview	/SPC1_RESULTS
SPC1_METRICS_0_Raw_Results.xlsx	Raw results for Primary Metrics Test	/SPC1_RESULTS
SPC1_METRICS_0_Summary_Results.xlsx	Primary Metrics Summary	/SPC1_RESULTS
SPC1_PERSIST_1_0_Raw_Results.xlsx	Raw results for PERSIST1 Test Phase	/SPC1_RESULTS
SPC1_PERSIST_2_0_Raw_Results.xlsx	Raw results for PERSIST2 Test Phase	/SPC1_RESULTS
SPC1_Run_Set_Overview.xlsx	Run Set Overview Worksheet	/SPC1_RESULTS
SPC1_VERIFY_0_Raw_Results.xlsx	Raw results for first VERIFY Test Phase	/SPC1_RESULTS
SPC1_VERIFY_1_Raw_Results.xlsx	Raw results for second VERIFY Test Phase	/SPC1_RESULTS
/C_Tuning	**Tuning parameters and options**	**root**
ai-max-nr.sh	Adjust aio-max-nr	/C_Tuning
/D_Creation	**Storage configuration creation**	**root**
auto_create.py	Create storage cluster	/D_Creation
connect_nvme.sh	Map all LUNs on hosts as NVMe disk	/D_Creation
init_wds.sh	format storage cluster	/D_Creation
lvm.sh	Create logical volumes	/D_Creation
lvm_active.sh	activate logical volumes	/D_Creation
/E_Inventory	**Configuration inventory**	**root**
profile_WDS.sh	Captures profile of storage environment	/E_Inventory
volume_list.sh	Captures list of logical volumes	/E_Inventory
profile_WDS-init.txt	Storage configuration before INIT	/E_Inventory
profile_WDS-verify1.txt	Storage configuration before VERIFY1	/E_Inventory
profile_WDS-metrics.txt	Storage configuration before METRICS	/E_Inventory
profile_WDS-verify2.txt	Storage configuration before VERIFY2	/E_Inventory
profile_WDS-persist1.txt	Storage configuration before PERSIST1	/E_Inventory
profile_WDS-persist2.txt	Storage configuration before PERSIST2	/E_Inventory
profile_WDS-end.txt	Storage configuration after test	/E_Inventory
volume_list_init.txt	List of logical volumes before INIT	/E_Inventory
volume_list_verify1.txt	List of logical volumes before VERIFY1	/E_Inventory
volume_list_metrics.txt	List of logical volumes before METRICS	/E_Inventory

volume_list_verify2.txt	List of logical volumes before VERIFY2	/E_Inventory
volume_list_persist1.txt	List of logical volumes before PERSIST1	/E_Inventory
volume_list_persist2.txt	List of logical volumes before PERSIST2	/E_Inventory
volume_list_end.txt	List of logical volumes after test	/E_Inventory
dmidecode_info.sh	Collect node system info	/E_Inventory
dmidecode_info_init.txt	List node system info	/E_Inventory
lsblk_info.sh	Collect node disk info	/E_Inventory
lsblk_info_init.txt	List node disk info	/E_Inventory
network_info.sh	Collect node network info	/E_Inventory
network_info_init.txt	List node network info	/E_Inventory
Operate_System_info.sh	Collect node OS info	/E_Inventory
OS_info_init.txt	List node OS info	/E_Inventory
/F_Generator	**Workload generator**	**root**
HOST32.HST	Host configuration file	/F_Generator
slave_asu.asu	Define LUNs hosting the ASUs	/F_Generator
WDS_init.sh	Execute INIT	/F_Generator
WDS_metrics.sh	Execute METRICS	/F_Generator
WDS_persist1.sh	Execute PERSIST1	/F_Generator
WDS_persist2.sh	Execute PERSIST2	/F_Generator
WDS_verify1.sh	Execute VERIFY1	/F_Generator
WDS_verify2.sh	Execute VERIFY1	/F_Generator

APPENDIX B: THIRD PARTY QUOTATION

All components are available directly through the Test Sponsor (Huarui Expon Technologies).

APPENDIX C: TUNING PARAMETERS AND OPTIONS

See Appendix D Step 4.

APPENDIX D: STORAGE CONFIGURATION CREATION

Step 1: Create Storage Pools, USS and NVMe-oF

Execute init_wds.sh to complete the following:

- Create 8 storage pools
- Create 4 USS gateways on each host node, 128 USS total
- Create 48 LUNs (6 LUN per Pool, 558 GiB per LUN)
- Create 4 NVMe-oF
- Map LUNs and USS to the 4 NVMe-oF

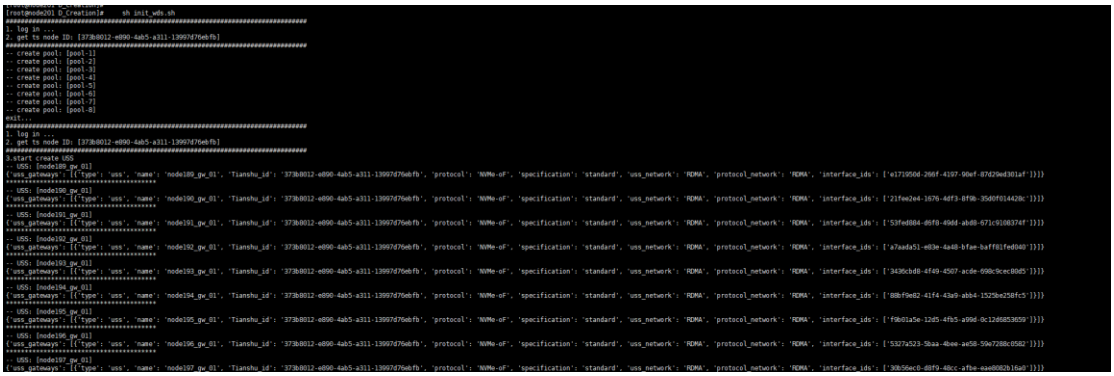

Step 2: map LUNs as NVMe disks on host nodes

Run connect_nvme.sh script to map all LUNs on hosts as NVMe disk.

```
[root@node201 D_Creation]# date
Thu Aug 31 17:12:34 CST 2023
[root@node201 D_Creation]# sh connect_nvme.sh
------------------------------ 10.1.100.189: ------------------------------
start discover
=====Discovery Log Entry 0======
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not required
portid:  0
trsvcid: 4421
subnqn:  nqn.2023-08.com.sds.wds:nvmf01
traddr:  10.1.100.189
rdma_prtype: not specified
rdma_qptype: connected
rdma_cms:    rdma-cm
rdma_pkey: 0x0000
=====Discovery Log Entry 0======
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not required
portid:  0
trsvcid: 4422
subnqn:  nqn.2023-08.com.sds.wds:nvmf02
traddr:  10.1.100.189
rdma_prtype: not specified
rdma_qptype: connected
rdma_cms:    rdma-cm
rdma_pkey: 0x0000
=====Discovery Log Entry 0======
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not required
portid:  0
trsvcid: 4423
```

Step 3: Create logical volumes, activate them on each host node

Execute the lvm.sh script to Create 48 Physical Volumes. Create 1 Volume Group(vg1) using 48 Physical Volumes. Create 18 Logical Volumes, each with a capacity of 640 GiB, on vg1 for ASU-1. Create 18 Logical Volumes, each with a capacity of 640 GiB, on vg1 for ASU-2. Create 4 Logical Volumes, each with a capacity of 640 GiB, on vg1 for ASU-3.

```
[root@node201 D_Creation]# sh lvm.sh
--------persist node:189
persist_nvme.sh
/dev/nvme0n1 uuid is dfa26a3a-4669-11ee-b165-e878ee3b64d0 nsid is 1 ,new udev disk is wds/nvme0n1
/dev/nvme0n10 uuid is dfebc51b-4669-11ee-b8c7-e878ee3b5d44 nsid is 10 ,new udev disk is wds/nvme0n10
/dev/nvme0n11 uuid is dff16f4a-4669-11ee-b165-e878ee3b64d0 nsid is 11 ,new udev disk is wds/nvme0n11
/dev/nvme0n12 uuid is dff7231e-4669-11ee-9f54-e878ee3b67c4 nsid is 12 ,new udev disk is wds/nvme0n12
/dev/nvme0n2 uuid is dfaaf950-4669-11ee-9f54-e878ee3b67c4 nsid is 2 ,new udev disk is wds/nvme0n2
/dev/nvme0n3 uuid is dfb3580d-4669-11ee-b165-e878ee3b64d0 nsid is 3 ,new udev disk is wds/nvme0n3
/dev/nvme0n4 uuid is dfbd24ac-4669-11ee-b8c7-e878ee3b5d44 nsid is 4 ,new udev disk is wds/nvme0n4
/dev/nvme0n5 uuid is dfc4fd5f-4669-11ee-9f54-e878ee3b67c4 nsid is 5 ,new udev disk is wds/nvme0n5
/dev/nvme0n6 uuid is dfccfc26-4669-11ee-b8c7-e878ee3b5d44 nsid is 6 ,new udev disk is wds/nvme0n6
/dev/nvme0n7 uuid is dfd497c8-4669-11ee-b165-e878ee3b64d0 nsid is 7 ,new udev disk is wds/nvme0n7
/dev/nvme0n8 uuid is dfdd1e7b-4669-11ee-9f54-e878ee3b67c4 nsid is 8 ,new udev disk is wds/nvme0n8
/dev/nvme0n9 uuid is dfe51625-4669-11ee-b8c7-e878ee3b5d44 nsid is 9 ,new udev disk is wds/nvme0n9
/dev/nvme1n1 uuid is dffdbb56-4669-11ee-b8c7-e878ee3b5d44 nsid is 1 ,new udev disk is wds/nvme1n1
/dev/nvme1n10 uuid is e03c6d80-4669-11ee-b8c7-e878ee3b5d44 nsid is 10 ,new udev disk is wds/nvme1n10
/dev/nvme1n11 uuid is e04338ac-4669-11ee-9f54-e878ee3b67c4 nsid is 11 ,new udev disk is wds/nvme1n11
/dev/nvme1n12 uuid is e04c60e1-4669-11ee-b165-e878ee3b64d0 nsid is 12 ,new udev disk is wds/nvme1n12
/dev/nvme1n2 uuid is e0042dd0-4669-11ee-9f54-e878ee3b67c4 nsid is 2 ,new udev disk is wds/nvme1n2
/dev/nvme1n3 uuid is e00b8450-4669-11ee-b8c7-e878ee3b5d44 nsid is 3 ,new udev disk is wds/nvme1n3
/dev/nvme1n4 uuid is e012bc12-4669-11ee-b165-e878ee3b64d0 nsid is 4 ,new udev disk is wds/nvme1n4
/dev/nvme1n5 uuid is e01a82b4-4669-11ee-b8c7-e878ee3b5d44 nsid is 5 ,new udev disk is wds/nvme1n5
/dev/nvme1n6 uuid is e023765e-4669-11ee-b165-e878ee3b64d0 nsid is 6 ,new udev disk is wds/nvme1n6
/dev/nvme1n7 uuid is e02ad7b7-4669-11ee-b8c7-e878ee3b5d44 nsid is 7 ,new udev disk is wds/nvme1n7
/dev/nvme1n8 uuid is e0306d9a-4669-11ee-b165-e878ee3b64d0 nsid is 8 ,new udev disk is wds/nvme1n8
/dev/nvme1n9 uuid is e035ea5c-4669-11ee-9f54-e878ee3b67c4 nsid is 9 ,new udev disk is wds/nvme1n9
/dev/nvme2n1 uuid is e0539262-4669-11ee-b165-e878ee3b64d0 nsid is 1 ,new udev disk is wds/nvme2n1
/dev/nvme2n10 uuid is e0a0502e-4669-11ee-9f54-e878ee3b67c4 nsid is 10 ,new udev disk is wds/nvme2n10
/dev/nvme2n11 uuid is e0a75606-4669-11ee-b165-e878ee3b64d0 nsid is 11 ,new udev disk is wds/nvme2n11
/dev/nvme2n12 uuid is e0ace0c1-4669-11ee-b8c7-e878ee3b5d44 nsid is 12 ,new udev disk is wds/nvme2n12
/dev/nvme2n2 uuid is e05d5002-4669-11ee-b8c7-e878ee3b5d44 nsid is 2 ,new udev disk is wds/nvme2n2
/dev/nvme2n3 uuid is e06825dd-4669-11ee-b8c7-e878ee3b5d44 nsid is 3 ,new udev disk is wds/nvme2n3
/dev/nvme2n4 uuid is e0726bf6-4669-11ee-b8c7-e878ee3b5d44 nsid is 4 ,new udev disk is wds/nvme2n4
/dev/nvme2n5 uuid is e0784d13-4669-11ee-b165-e878ee3b64d0 nsid is 5 ,new udev disk is wds/nvme2n5
/dev/nvme2n6 uuid is e0813550-4669-11ee-9f54-e878ee3b67c4 nsid is 6 ,new udev disk is wds/nvme2n6
/dev/nvme2n7 uuid is e08b92d9-4669-11ee-b8c7-e878ee3b5d44 nsid is 7 ,new udev disk is wds/nvme2n7
/dev/nvme2n8 uuid is e0941ef4-4669-11ee-b165-e878ee3b64d0 nsid is 8 ,new udev disk is wds/nvme2n8
/dev/nvme2n9 uuid is e09add0c-4669-11ee-b8c7-e878ee3b5d44 nsid is 9 ,new udev disk is wds/nvme2n9
/dev/nvme3n1 uuid is e0b3973b-4669-11ee-9f54-e878ee3b67c4 nsid is 1 ,new udev disk is wds/nvme3n1
/dev/nvme3n10 uuid is e0eb336f-4669-11ee-9f54-e878ee3b67c4 nsid is 10 ,new udev disk is wds/nvme3n10
/dev/nvme3n11 uuid is e0f0b6f9-4669-11ee-b8c7-e878ee3b5d44 nsid is 11 ,new udev disk is wds/nvme3n11
/dev/nvme3n12 uuid is e0f66654-4669-11ee-b165-e878ee3b64d0 nsid is 12 ,new udev disk is wds/nvme3n12
/dev/nvme3n2 uuid is e0b9781b-4669-11ee-b165-e878ee3b64d0 nsid is 2 ,new udev disk is wds/nvme3n2
/dev/nvme3n3 uuid is e0bf8db9-4669-11ee-9f54-e878ee3b67c4 nsid is 3 ,new udev disk is wds/nvme3n3
/dev/nvme3n4 uuid is e0c536ed-4669-11ee-b8c7-e878ee3b5d44 nsid is 4 ,new udev disk is wds/nvme3n4
/dev/nvme3n5 uuid is e0cbc994-4669-11ee-9f54-e878ee3b67c4 nsid is 5 ,new udev disk is wds/nvme3n5
/dev/nvme3n6 uuid is e0d2f203-4669-11ee-9f54-e878ee3b67c4 nsid is 6 ,new udev disk is wds/nvme3n6
/dev/nvme3n7 uuid is e0d97c2b-4669-11ee-9f54-e878ee3b67c4 nsid is 7 ,new udev disk is wds/nvme3n7
/dev/nvme3n8 uuid is e0df3afb-4669-11ee-b8c7-e878ee3b5d44 nsid is 8 ,new udev disk is wds/nvme3n8
/dev/nvme3n9 uuid is e0e5cd40-4669-11ee-b8c7-e878ee3b5d44 nsid is 9 ,new udev disk is wds/nvme3n9
--------end persist node:189
--------persist node:190
```

Execute the lvm_active.sh to make each logical volume available (activate)

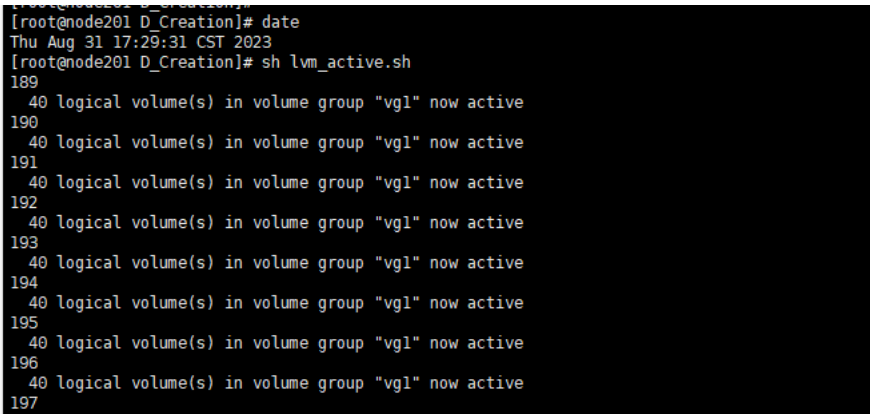

Step 4: Change the Scheduler on each Host System

Change the scheduler on each host system by executing the set_nr_requests.sh script on each host system. This will change the maximum number of AIO operations to 1048576.

APPENDIX E: CONFIGURATION INVENTORY

The scripts used to collect the configuration inventory and the log files that were generated are available in the Supporting Files (see Appendix_A).

APPENDIX F: WORKLOAD GENERATOR

The ASUs accessed by the SPC-1 workload generator are defined in slave_asu.asu. The hosts used to drive the SPC-1 workload are defined in HOST32.HST. The scripts used to execute the benchmark sequence are:

- WDS_init.sh
- WDS_verify1.sh
- WDS_metrics.sh
- WDS_verify2.sh
- WDS_persist1.sh
- WDS_persist2.sh

The files are included in the Supporting Files (see Appendix A).